Exhibit 3

🐦 **Transparency**   Reports ⌄   About   Resources   🔍

Rules Enforcement

About this report
Insights into how and when we enforce our policies, and reports of potential violations.
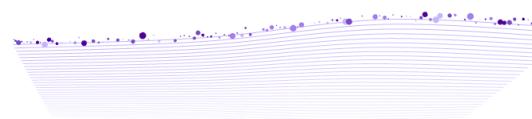
✕ **Rules Enforcement**   Accounts Actioned   Accounts Reported   **Jan - Jun 2020** ⌄   **Download Report** ⌄

Accounts Actioned

Published on January 11, 2021

01. **Latest Data**

02. **Overview**

03. **Analysis**

---

01. **Latest Data: Accounts Actioned**   (Bars Table) 🐦
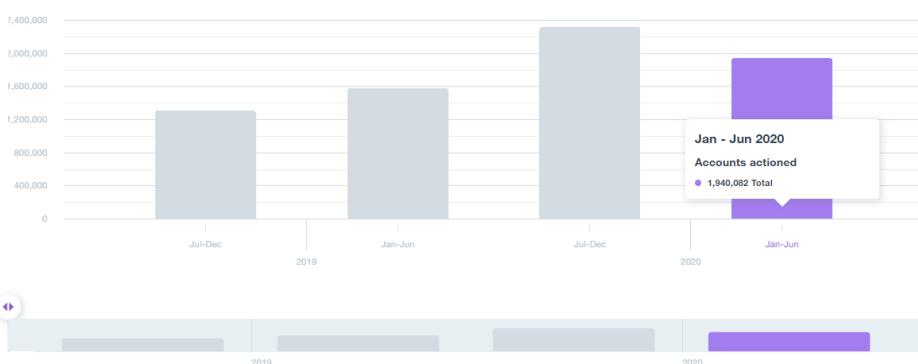
| Jan - Jun 2020 |
| Accounts actioned |
| • 1,940,082 Total |

Accounts actioned - January - June 2020

Accounts actioned	**Accounts suspended**	**Content removed**
1.9M	**925.7K**	**1.9M**

---

02. **Overview**

Twitter's purpose is to serve the public conversation. We welcome people to share their unique point of view on Twitter, but there are some behaviors that discourage others from expressing themselves or place people at risk of harm. The Twitter Rules exist to help ensure that all people can participate in the public conversation freely and safely, and include specific policies that explain the types of content and behavior that are prohibited.

This section covers the latest data about instances where we've taken enforcement actions under the Twitter Rules to either require the removal of specific Tweets or to suspend accounts. These metrics are referred to as: accounts actioned, content removed, and accounts suspended. More details about our range of enforcement options are available in our Help Center.

Twitter's operations were affected due to the unprecedented COVID-19 pandemic. Starting in March, the majority of our global operations centers were temporarily closed due to lockdown orders and related health concerns, significantly

reducing our human review capacity. We maintained dedicated resources focused on reviewing and taking enforcement action against content most likely to cause severe harm (for example, child sexual exploitation and terrorism) in addition to prioritizing reports where we were able to predict a high likelihood of a Rules violation. As a result of this prioritization, we saw significant slowdowns and backlogs in other areas, and have continued to evolve our approach.

Some notable changes since our last report:

Accounts actioned

-16%

Decrease in **accounts actioned** compared to the last reporting period.

Accounts suspended

+6%

Increase in **accounts suspended** compared to the last reporting period.

Content removed

-33%

Decrease in **content removed** compared to the last reporting period.
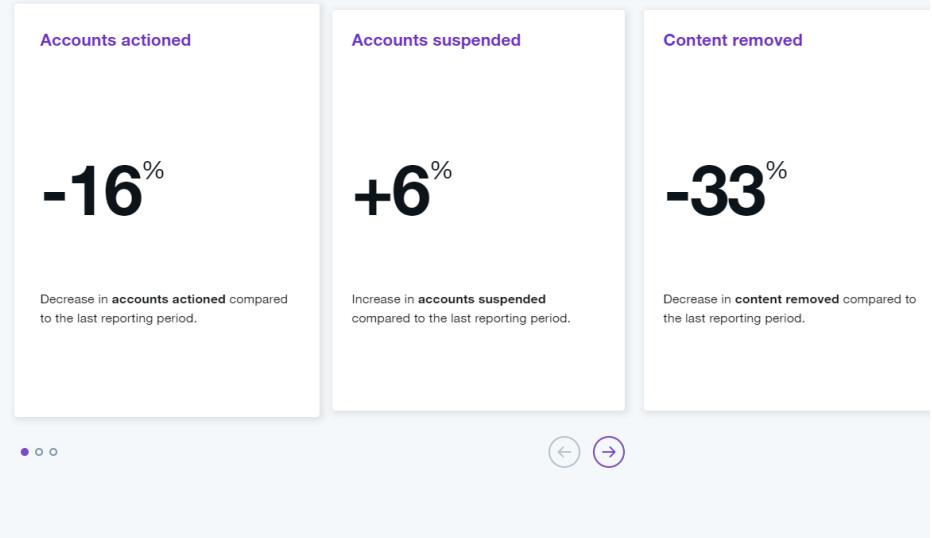
03. Analysis

Bars | Table

Accounts actioned - January - June 2020

Accounts actioned	Accounts suspended	Content removed
1.9M	925.7K	1.9M
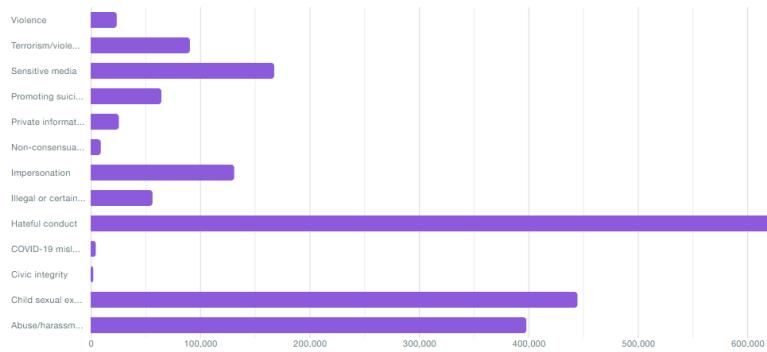
Big picture

We have a global team that manages enforcement of the Twitter Rules with 24/7 coverage in every supported language on Twitter. Our goal is to apply the Twitter Rules objectively and consistently. Enforcement actions are taken on content that is determined to violate the Twitter Rules.

We support the spirit of the [Santa Clara Principles on Transparency and Accountability in Content Moderation](#), and are committed to sharing more detailed information about how we enforce the Twitter Rules in future reports.

Safety

The "Safety" section of the Twitter Rules covers violence, terrorism/violent extremism, child sexual exploitation, abuse/harassment, hateful conduct, promoting suicide or self-harm, sensitive media (including graphic violence and adult content), and illegal or certain regulated goods or services. More information about each policy can be found in the [Twitter Rules](#).

Some notable changes since the last report:

Violence

There was a **48% decrease** in the number of accounts actioned for violations of our violence policies during this reporting period.

Terrorism/violent extremism

There was a **5% increase** in the number of accounts actioned for violations of our terrorism / violent extremism policy during this reporting period.

Child sexual exploitation

There was a **68% increase** in the number of accounts actioned for violations of our child sexual exploitation policy during this reporting period.

Other select takeaways:

Terrorism/violent extremism
The Twitter Rules prohibit the promotion of terrorism and violent extremism. Action was taken on 90,684 unique accounts under this policy during this reporting period. 94% of those accounts were proactively identified and actioned. Our current methods of surfacing potentially violating content for review include leveraging the shared industry hash database supported by the Global Internet Forum to Counter Terrorism (GIFCT).

Child sexual exploitation
We do not tolerate child sexual exploitation on Twitter. When we are made aware of child sexual exploitation media, including links to images of or content promoting child exploitation, the material will be removed from the site without further notice and reported to The National Center for Missing & Exploited Children ("NCMEC"). People can report content that appears to violate the Twitter Rules regarding Child Sexual Exploitation via our web form or through in-app reporting.

438,809 unique accounts were suspended during this reporting period for violating Twitter policies prohibiting child sexual exploitation. 91% of those accounts were proactively identified by employing internal proprietary tools and industry hash sharing initiatives. These tools and initiatives support our efforts to surface potentially violative content for further review and, if appropriate, removal.

Sensitive media, including graphic violence and adult content
These policies saw the largest increase in the number of accounts actioned during this reporting period.

Hateful conduct
Hateful conduct expanded to include a new dehumanization policy on March 5, 2020.

Privacy

The "Privacy" section of the Twitter Rules covers private information and non-consensual nudity. More information about each policy can be found in the Twitter Rules.

Some notable changes since the last report:

Private information

There was a **14% increase** in the number of accounts actioned for violations of our private information policy during this reporting period.

Non-consensual nudity

There was a **48% decrease** in the number of accounts actioned for violations of our non-consensual nudity policy during this reporting period.

Other select takeaways:

Private information

This reporting period saw the largest increase in the number of accounts actioned under this policy. Internal tooling improvements allowed us to increase enforcement of this policy.

Authenticity

The "Authenticity" section of the Twitter Rules covers platform manipulation and spam, civic integrity, impersonation, synthetic and manipulated media, and copyright and trademark. We have standalone report pages for platform manipulation and spam, copyright, and trademark, and cover civic integrity and impersonation enforcement actions in this section.[2] More information about each policy can be found in the Twitter Rules.

Some notable changes since the last report:

Civic integrity

There was a **37% increase** in the number of accounts actioned for violations of our civic integrity policy during this reporting period.

Impersonation

There was a **28% decrease** in the number of accounts actioned for violations of our impersonation policy during this reporting period.

COVID-19 misleading misinformation

We **suspended or required the removal of content from 4,658 accounts** for violations of our COVID-19 misleading information policy during this reporting period. This number does not include accounts where we applied a label or warning message.

Other select takeaways:

Civic Integrity

This reporting period saw an increase in the number of accounts actioned under this policy. Enforcements increased in the lead up to the US elections in November 2020.

Accounts Reported

Published on December 18, 2020

01.   **Latest Data**

02.   **Overview**

03.   **Analysis**

01.   Latest Data: Accounts Reported

Bars   Table
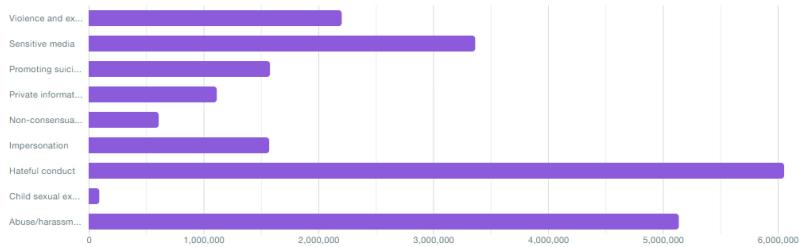
Accounts reported - January - June 2020

Accounts reported

12.4M

02. **Overview**

Insights into <u>accounts reported</u> for violations of the Twitter Rules.

Accounts reported

+30[%]

Increase in **accounts reported** compared to the last reporting period.
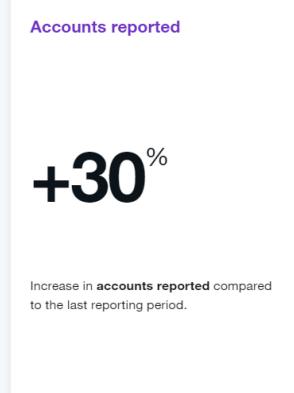
03. **Analysis**

Big picture

Reported content is reviewed to determine whether it violates any aspects of the Twitter Rules, *independent* of its initial report category. For example, content reported under our private information policy may be found to violate – and be actioned under – our hateful conduct policies. We may also determine that reported content does not violate the Rules at all.

The policy categories in this section do not map cleanly to the ones in the Accounts Actioned section above. This is because people typically report content for possible Twitter Rules violations through our Help Center or in-app reporting.

We support the spirit of the Santa Clara Principles on Transparency and Accountability in Content Moderation, and are committed to sharing more detailed information about how we enforce the Twitter Rules in future reports.

Footnotes

Accounts Actioned

To provide meaningful metrics, we de-duplicate accounts which were actioned multiple times for the same policy violation. This means that if we took action on a Tweet or account under multiple policies, the account would be counted separately under each policy. However, if we took action on a Tweet or account multiple times under the same policy (for example, we may have placed an account in read-only mode temporarily and then later also required media or profile edits on the basis of the same violation), the account would be counted once under the relevant policy.

1. Our synthetic and manipulated media policy launched in February 2020 and, as such, there is no enforcement data to share for this reporting period. We plan to include this information in future reports.

Accounts Reported

To provide meaningful metrics, we de-duplicate accounts which were reported multiple times (whether multiple users reported an account for the same potential violation, or whether multiple users reported the same account for different potential violations). For the purposes of these metrics, we similarly de-duplicate reports of specific Tweets. This means that even if we received reports about multiple Tweets by a single account, we only counted these reports towards the "accounts reported" metric once.

Other reports

Information Requests

Platform Manipulation

Information Operations

Legal requests for account information

Malicious automation and spam

Disclosures and elections integrity